

> Step 1

The approach to solving this problem is relatively simple, once parameters of a bioinformatics simulation are understood. Simulations tend to run days or months. Accordingly, losing simulation data or having a system failure during simulation are catastrophic events. Availability is therefore a critical evaluation parameter. Additionally, the disk array will be accessed by 1000 parallel processors. Throughput will be a major concern.

> Step 2

The primary role of the power constraint in this problem is to prevent simply maximizing the all parameters in the disk array. Adding additional disks and controllers without justification will increase power consumption unnecessarily.

> Step 1

Consider the system must provide both backup and archiving. Accordingly, we will need multiple copies of our data and may be required to move those copies offsite. This makes none of the solutions optimal.

> Step 2

RAID or a second backup array provides high speed backup, but does not provide archival capabilities. Magnetic tape allows archiving, but can be exceptionally slow when comparing to disk backups. Online backup automatically achieves archiving, but can be even slower than disks.

> Step 1

The benchmarks must be evaluating backup throughput. Most other parameters that manage selection of a system are relatively well understood portability and the cost being the main issues to be evaluated.